



Computing at CDF



Mark Neubauer

Massachusetts Institute of Technology
for the CDF Collaboration

- **Introduction**
- **Computing requirements**
- **Central Analysis Facility**
- **Data Handling**
- **Toward the Grid**
- **Conclusions**





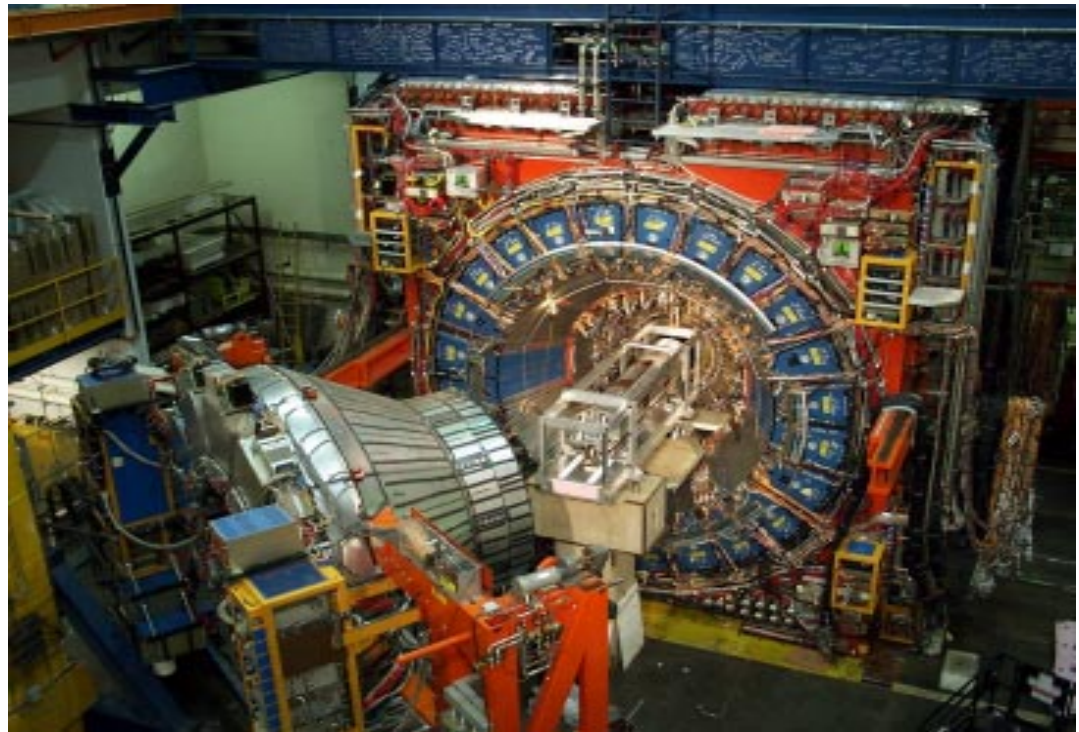
CDF in a Nutshell

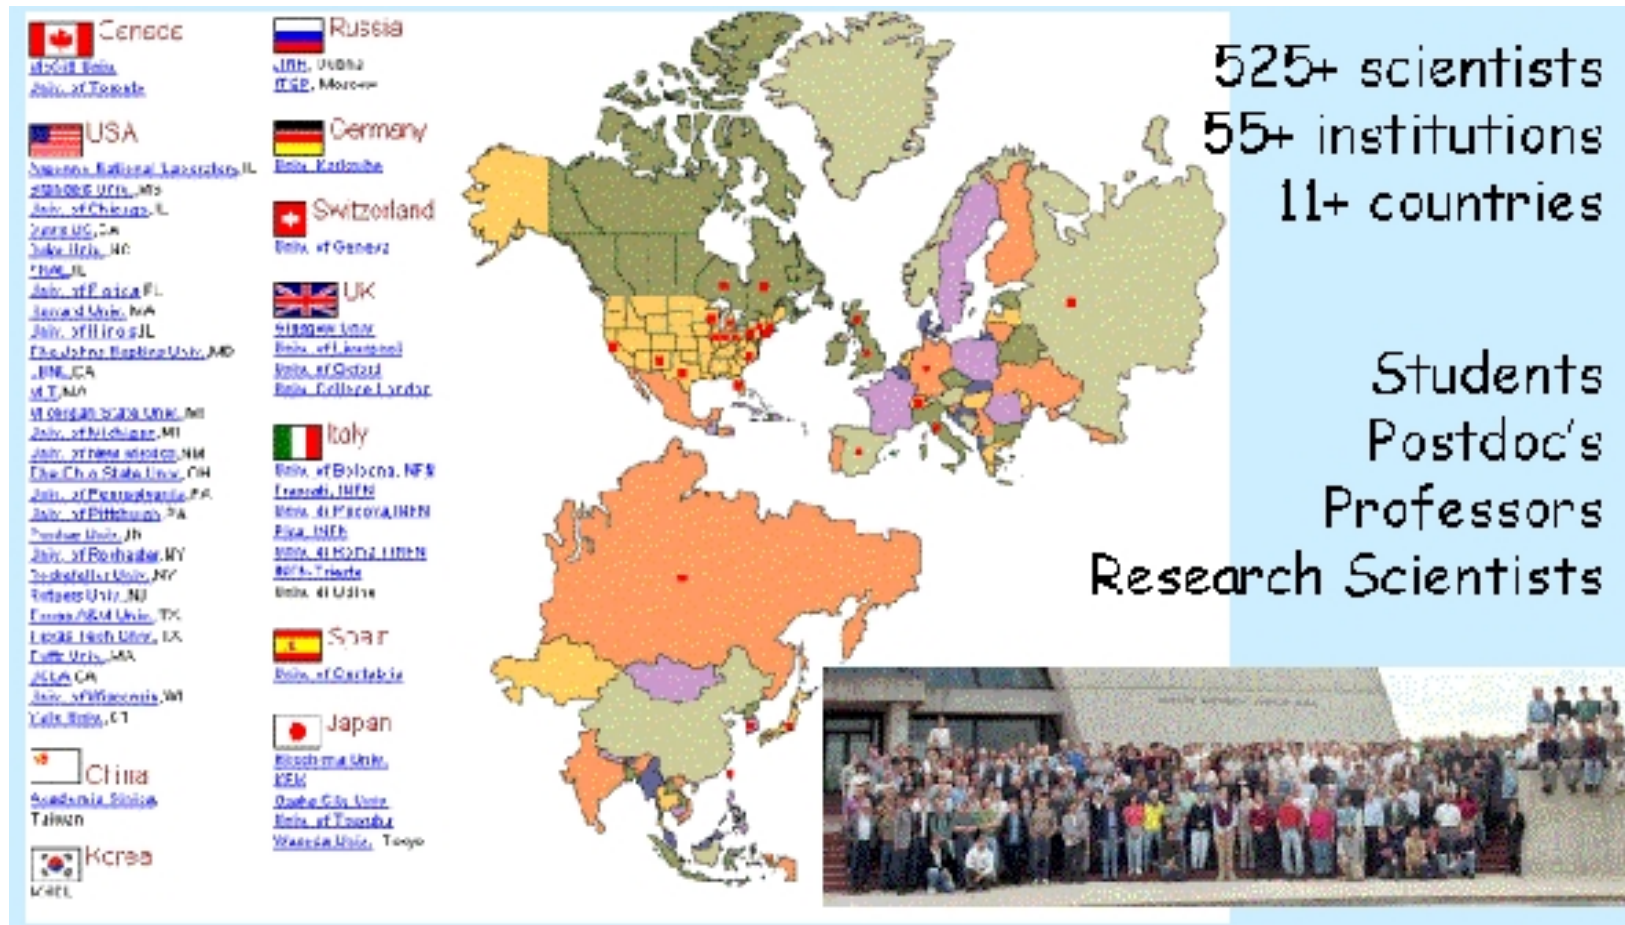


- CDF + D0 experiments analyze $p\bar{p}$ collisions from Tevatron at Fermilab
- Tevatron highest energy collider in world ($\sqrt{s} = 2 \text{ TeV}$) until LHC
- Run I (1992-1996) huge success \rightarrow 200+ papers (t quark discovery, ...)
- Run II (March 2001-) upgrades for luminosity ($\times 10$) + energy ($\sim 10\% \uparrow$)
 \rightarrow expect integrated luminosity $20\times$ (Run IIa) and $150\times$ (Run IIb) of Run I

Run II physics goals:

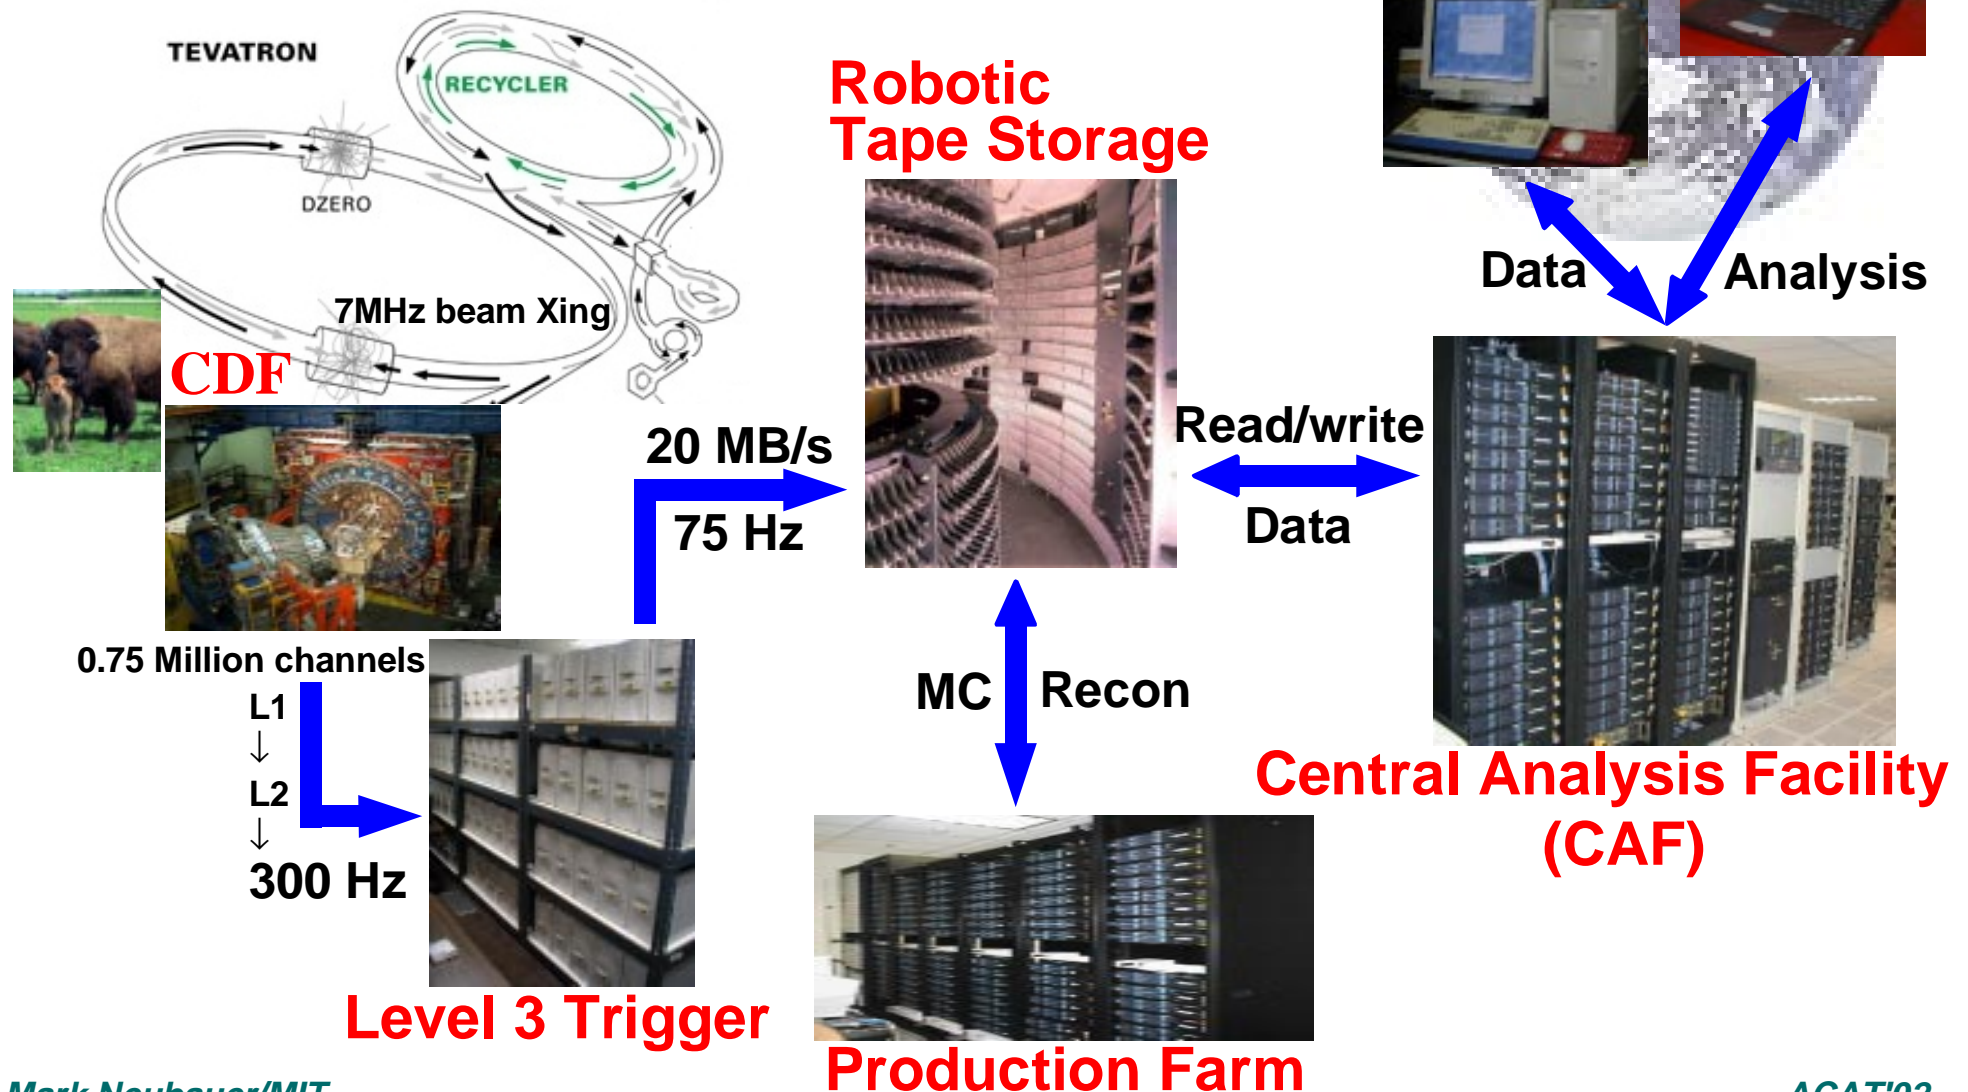
- **Search for Higgs boson**
- **Top quark properties** ($m_t, \sigma_{\text{tot}}, \dots$)
- **Electroweak** ($m_W, \Gamma_W, ZZ\gamma, \dots$)
- **Search for new physics** (e.g. SUSY)
- **QCD at large Q^2** (jets, α_s, \dots)
- **CKM tests in b hadron decays**





Goal: Provide computing resources for 200+ collaborators simultaneously doing analysis per day!

CDF DAQ/Analysis Flow





Reconstruction Farms

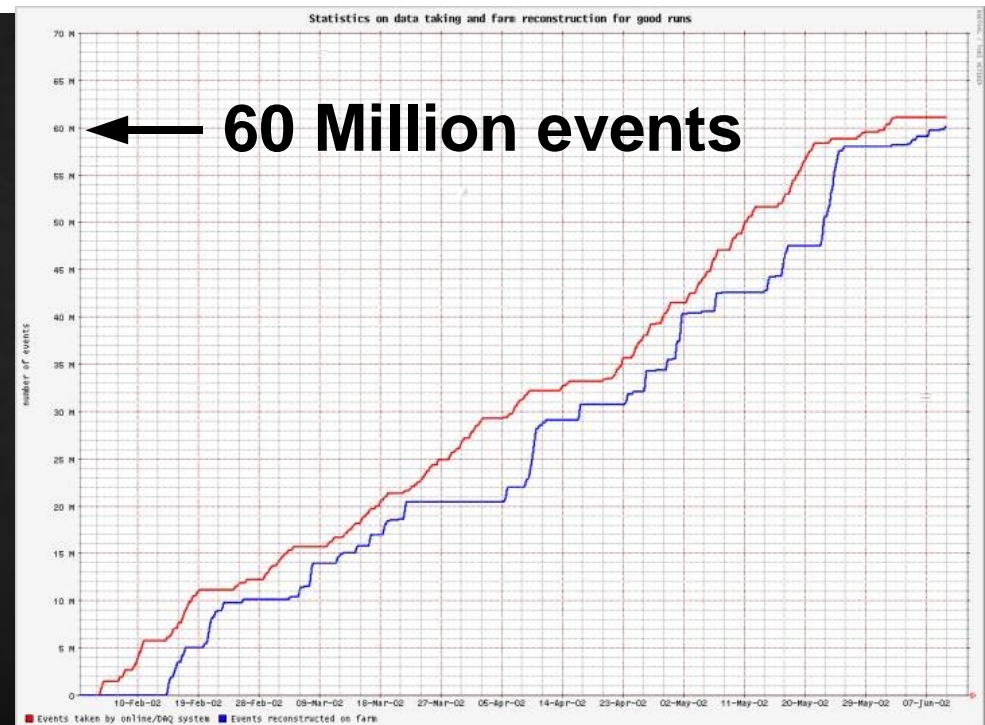


Data reconstruction + validation, Monte Carlo generation

154 dual P3's (equivalent to 244 1 Ghz machines)

Job management:

- Batch system → FBSNG developed at FNAL
- Single executable, validated offline





Database Usage at CDF



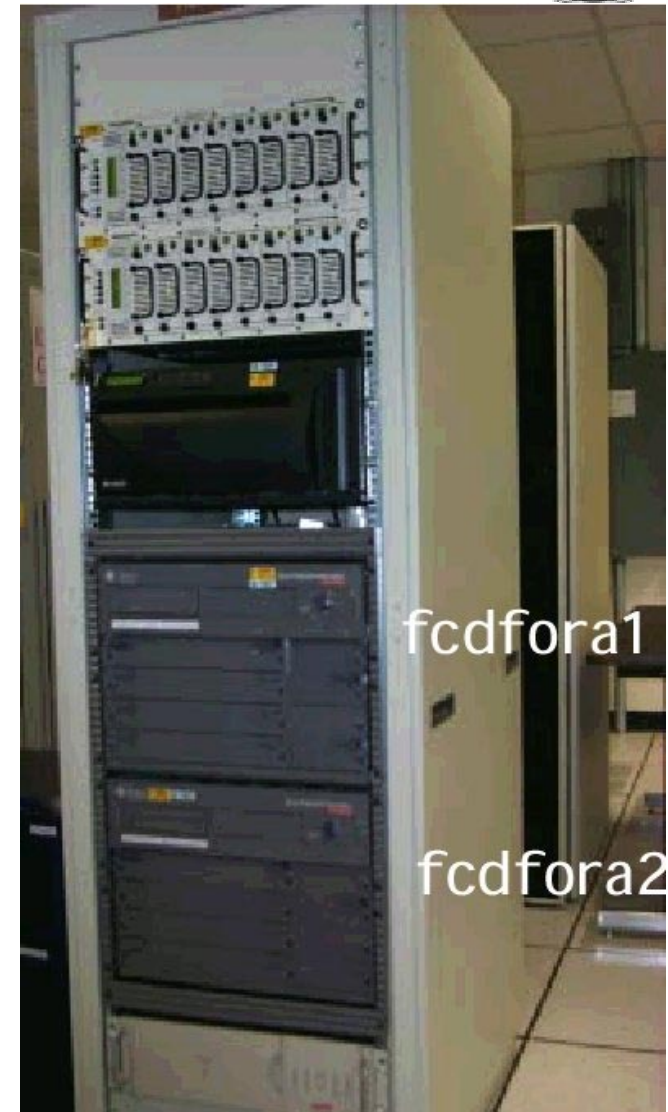
Oracle DB: Metadata + Calibrations

DB Hardware:

- **2 Sun E4500**

Presently evaluating:

- **Oracle on Linux**
- **MySQL**
- **Replication to remote sites**





Data/Software Characteristics



Data Characteristics:

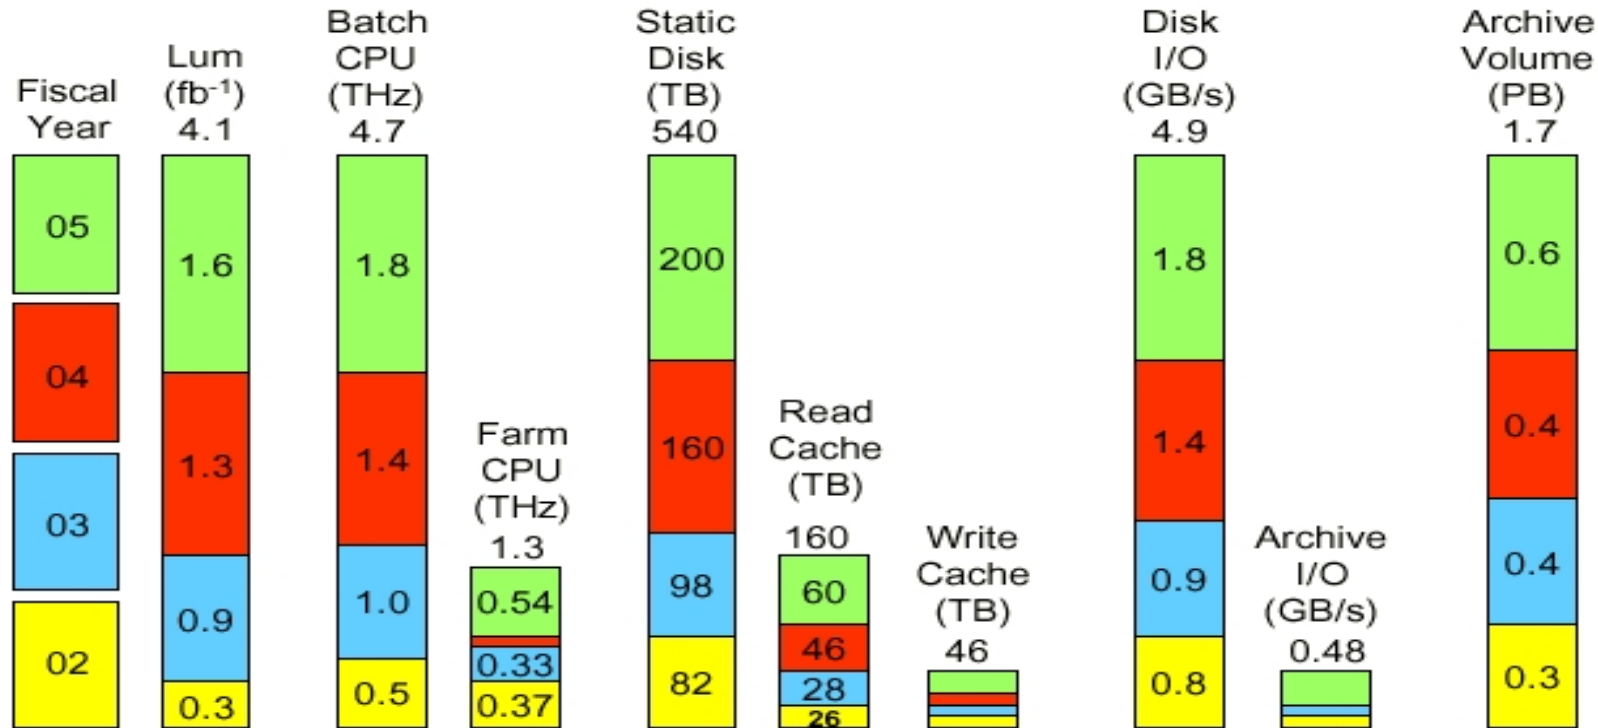
- Root I/O as persistent (raw) data format
- Raw data size: **~250 kB/event**
- Reconstructed data (PAD) format: **50-100 kB/event**
- Typical ntuple size (stntuple): **5-10 kB/event**
- Typical RunIIa secondary dataset size: **10^7 events**

Analysis Software:

- Typical analysis jobs run @ **10 Hz** on 1 GHz P3
→ **few MB/sec**
- CPU rather than I/O bound (FastEthernet)



Computing Requirements



Requirements set by goal:

200 simultaneous users to analyze secondary data set (10^7 evts) in a day

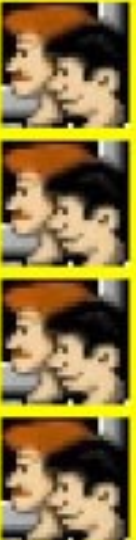
Need ~700 TB of disk and ~5 THz of CPU by end of FY'05:

→ need lots of disk → need cheap disk → IDE Raid

→ need lots of CPU → commodity CPU → dual Intel/AMD



Past CAF Computing Model



Large SMP (128 processor SGI)
Expensive disks (FiberChannel/SCSI)

Analysis Code Development
Analysis Job Debugging
Interactive Analysis Jobs
Batch Jobs
"Other" Usage

Very expensive to expand and maintain

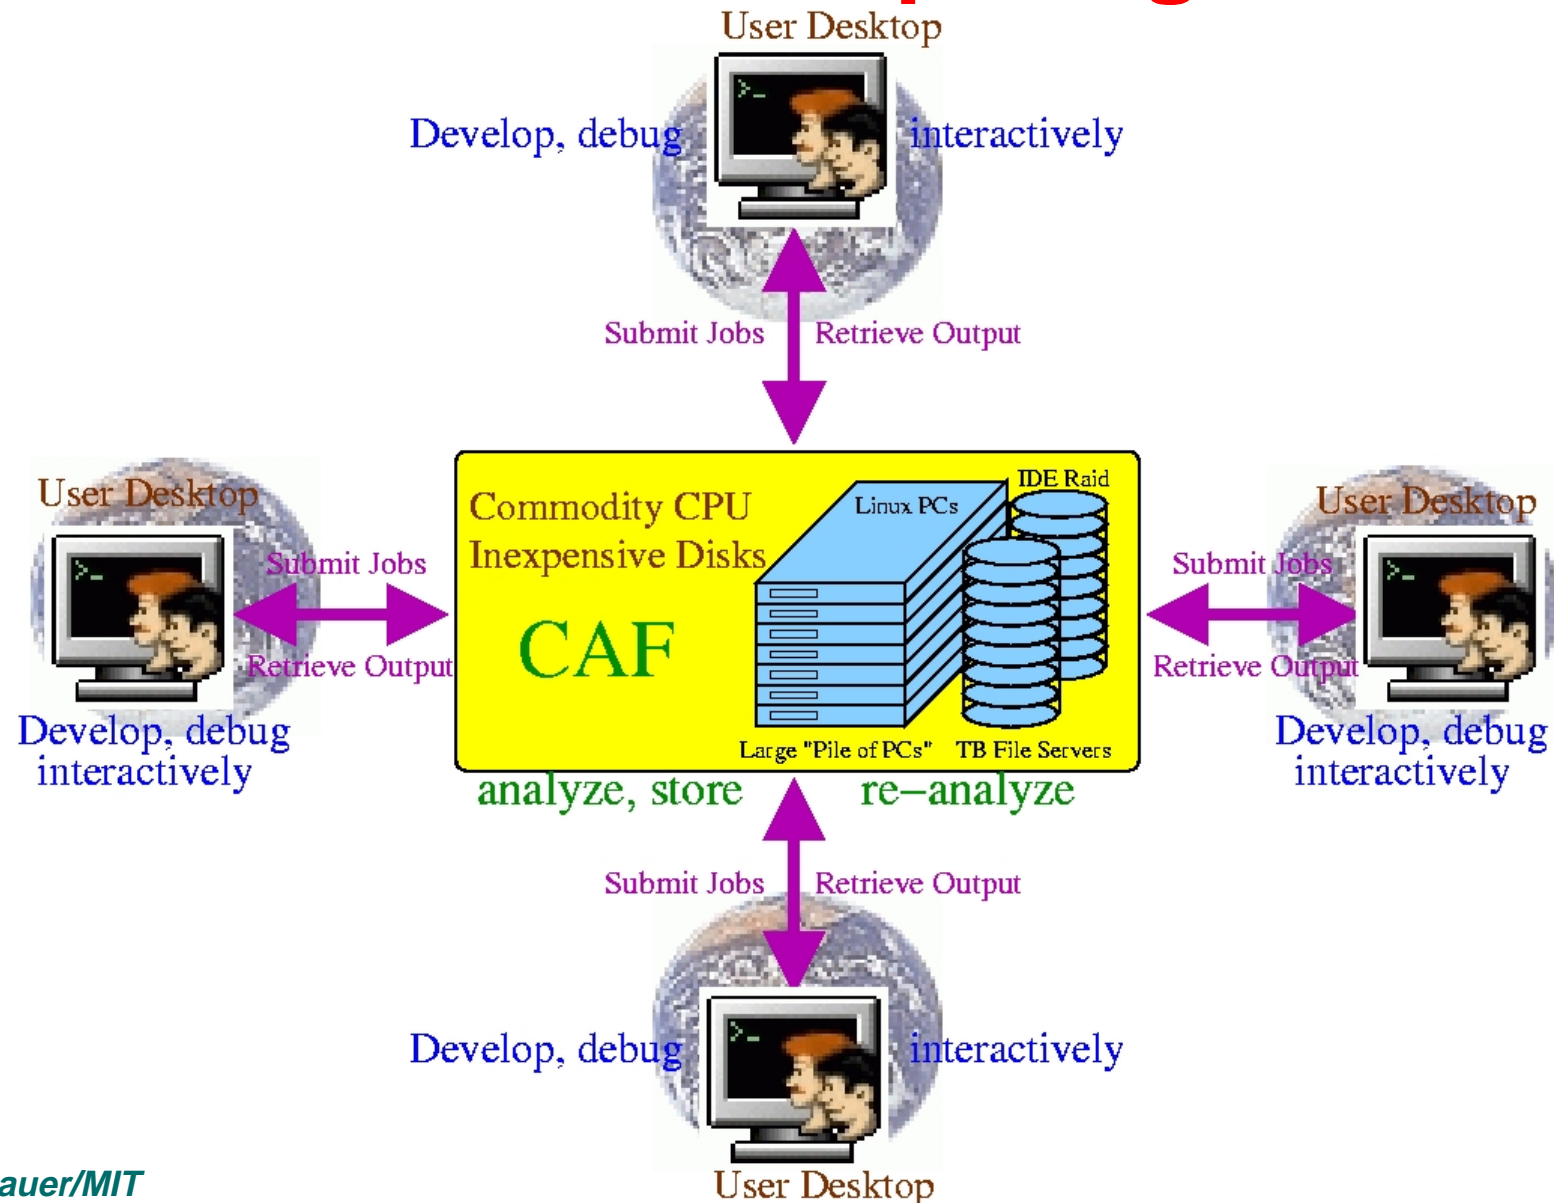
Bottom line:

Not enough 'bang for the buck'



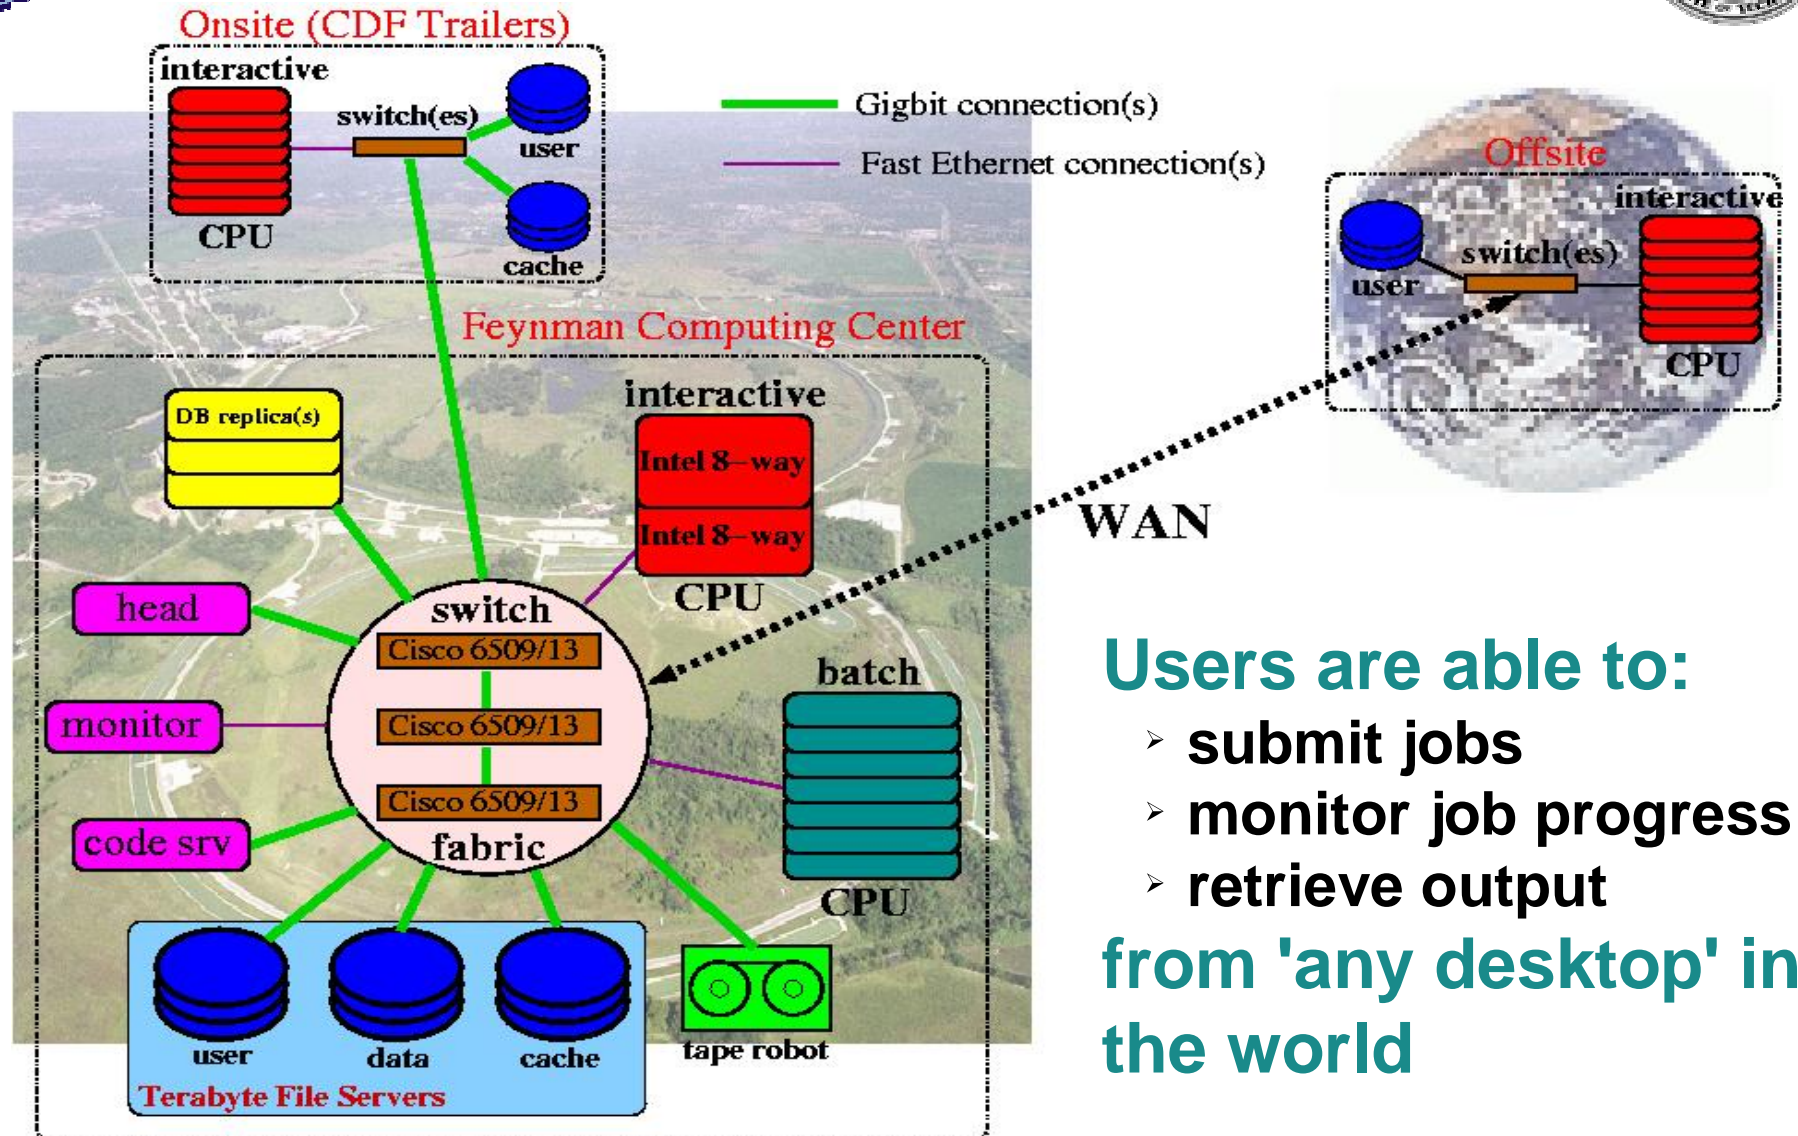


Present CAF Computing Model





CAF Implementation



Users are able to:

- submit jobs
- monitor job progress
- retrieve output

from 'any desktop' in the world



CAF Milestones



- Start of CAF design 11/01
- CAF prototype (protoCAF) assembled 2/25/02
- Fully-functional prototype system (>99% job success) 3/6/02
- ProtoCAF integrated into Stage1 system 4/25/02
- Production Stage1 CAF for collaboration 5/30/02



ProtoCAF



Stage1



Design → Production system in 6 months!



CAF Stage 1 Hardware



Code Server

File Servers

Worker Nodes

Linux 8-ways
(interactive)



Stage 1 Hardware: Workers



Workers (132 CPUs, 1U+2U rackmount):

16 2U Dual Athelon 1.6GHz / 512MB RAM

50 1U/2U Dual P3 1.26GHz / 2GB RAM

FE (11 MB/s) / 80GB job scratch each





Stage 1 Hardware: Servers



Mark Neubauer/MIT

Servers (35TB total, 16 4U rackmount):

2.2TB useable IDE RAID50 hot-swap

Dual P3 1.4GHz / 2GB RAM

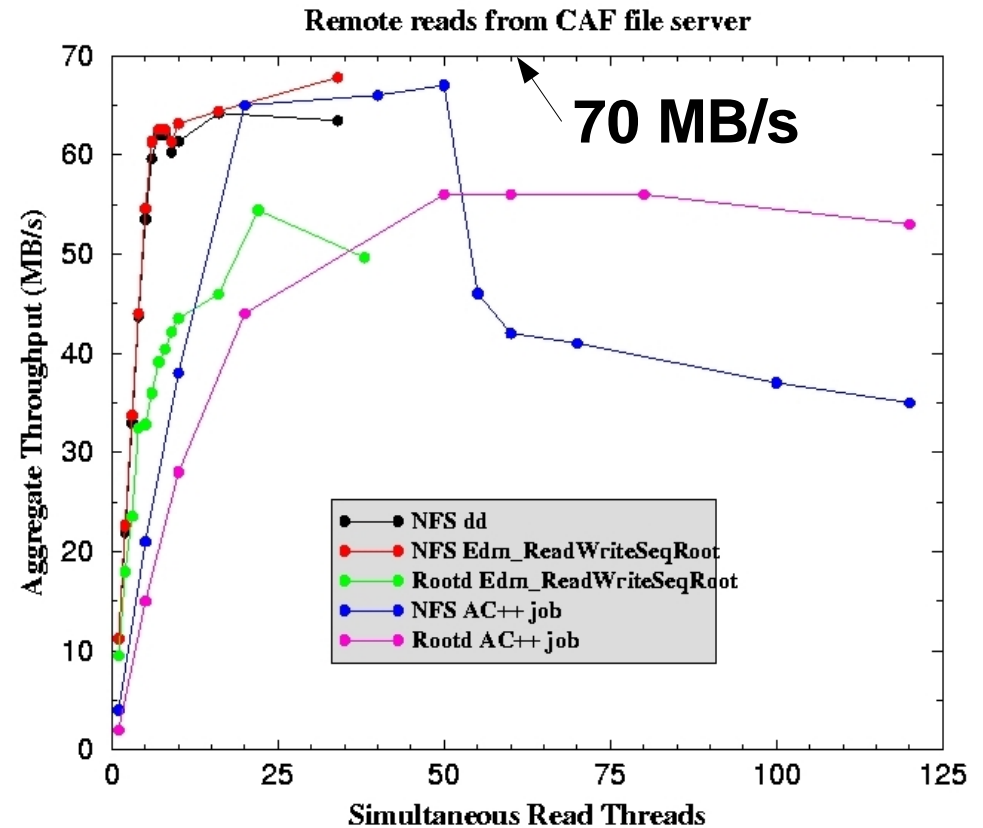
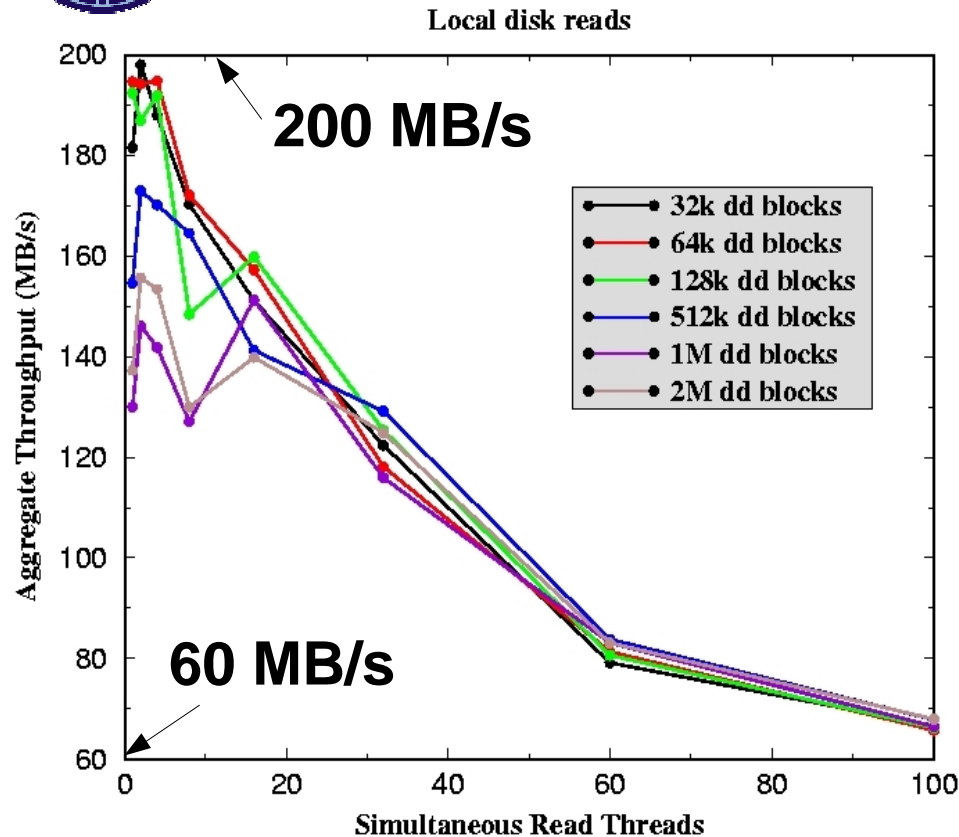
SysKonnct 9843 Gigabt Ethernet card



ACAT'02



File Server Performance



Server/Client Performance: Up to **200MB/s local reads**, **70 MB/s NFS**

Data Integrity tests: md5sum of local reads under heavy load

BER < 2×10^{-14} (Maxtor claims < 1 error / 10^{14} bits read)

Cooling tests: Temp profile of disks w/ IR gun after extended disk thrashing



CAF Software



Design goal:

Give users access to CAF resources

- **CPU**
- **scratch disk**
- **data handling system**

from their desktops anywhere in the world

Design constraints/desirables:

Fermilab computing security policy → **kerberos!**

Job scheduling → **proven batch system, configurable,
fair share capability, local support** → **FBSNG (FNAL-CD)**

Administrative ease → **no user accounts**

→ **non-interactive batch, jobs run as single 'cafuser' user**

User identity → **unique privileges for batch jobs, disk space**



User Access to CAF



Job Related:

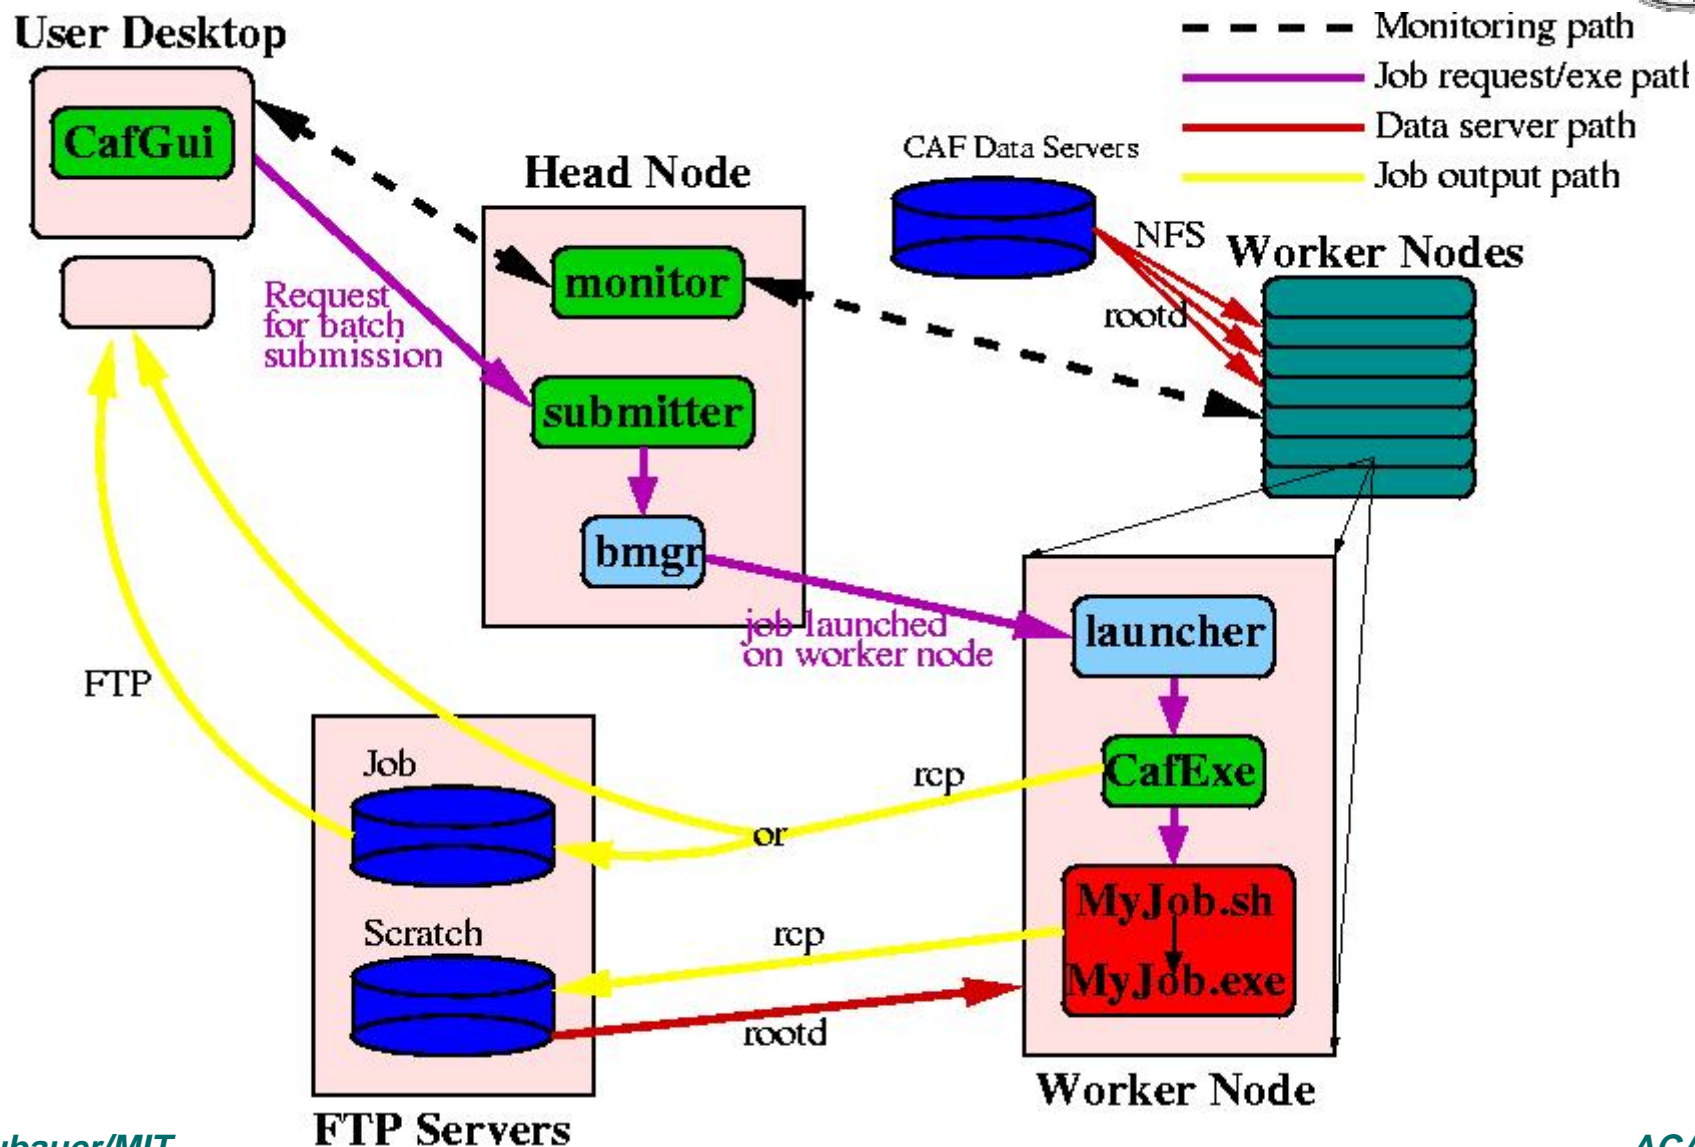
- **Submit jobs**
- **Check progress of job**
- **Kill a job**

Remote file system access:

- **'ls' in job's 'relative path'**
- **'ls' in a CAF node's absolute path**
- **'tail' of any file in job's 'relative path'**
- **Get full file listing based on metadata**



CAF Software





CAF User Interface



- Compile, build, debug analysis job on 'desktop'

- Fill in appropriate fields & submit job

section integer range

Initial Command: `/simple.sh`

Process Type: `Short`

Original Directory: `/home/msn/releases/development/CafUtil/examples`

Output File Location: `msn@fcdflnx2.fnal.gov/cdf/scratch/msn/temp/tgz`

☒ Email? Email Address: `msn@fnal.gov`

Ready

```
(2002-05-23 01:46:51) Email sent to msn@fnal.gov upon job completion
(2002-05-23 01:46:55) /bin/tar -cvzf /home/msn/msn49959.tgz *
(2002-05-23 01:46:57) Remove /home/msn/msn49959.tgz
(2002-05-23 01:46:57) Job Submission is successful, JID: 873
```

output destination

user exe+tcl directory

- Retrieve output using kerberized FTP tools
... or write output directly to 'desktop'!

Web Monitoring of User Queues

Each user a different queue

Process type for job length

test: 5 mins

short: 2 hrs

medium: 6 hrs

long: 2 days

This example:

1 job → 11 sections

(+ 1 additional section automatic for job cleanup)

| Name | Status | Default Process Type | Share | Prio | Waiting | Ready | Running | Total |
|----------|--------|----------------------|-------|------|---------|-------|---------|-------|
| akorn | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| amitl | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| anikeev | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| belforte | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| msmartin | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| msn | OK | short | 1.00 | 0 | 1 | 0 | 11 | 12 |
| pauly | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| paus | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| ratnikov | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| rescigno | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| semeria | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| sfiligoi | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| sgromoll | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| shepard | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| sidoti | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| spezziga | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| test | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| thkim | OK | short | 1.00 | 0 | 0 | 0 | 0 | 0 |
| thom | OK | short | 1.00 | 0 | 1 | 0 | 1 | 2 |

Monitoring jobs in your queue

The screenshot shows a Netscape browser window titled "Netscape: FBSWWW - queue msn@CAF". The main content area displays the "FBSNG on the web" interface. On the left, there is a "User Monitor" section. The main area shows a "Jobs" tab selected, displaying a table of jobs. The table has columns: SectID, User, ProcType, Status, Prio, NProc, and Date/Time. The jobs listed are 873.msn_600 through 873.msn_610, all with status "running", and 873.msn_end with status "waiting".

Queue Parameters [show]
Status: OK Running: 11 Pending: 0

| SectID | User | ProcType | Status | Prio | NProc | Date/Time |
|-------------|--------|----------|---------|------|-------|-----------------------------|
| 873.msn_600 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:09 |
| 873.msn_601 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:09 |
| 873.msn_602 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:10 |
| 873.msn_603 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:10 |
| 873.msn_604 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:11 |
| 873.msn_605 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:11 |
| 873.msn_606 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:12 |
| 873.msn_607 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:12 |
| 873.msn_608 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:12 |
| 873.msn_609 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:13 |
| 873.msn_610 | cdfcac | short | running | 0 | 1/1 | Started at 05/23 01:47:13 |
| 873.msn_end | cdfcac | mailer | waiting | 0 | 0/1 | Submitted at 05/23 01:46:57 |

FCS Group | FBSNG
FBSWWW version 0.1

Monitoring sections of your job

FBSNG on the web
Farm: CAF
Time: Thu May 23 01:47:23 2002
Report: Queue msn

Queues Jobs Nodes Process Types

Queue Parameters [show]

Status: OK Running: 11 Pending: 0

User Monitor

| SectID | User | ProcType |
|-------------|---------|----------|
| 873.msn_601 | cdcfcaf | short |
| 873.msn_602 | cdcfcaf | short |
| 873.msn_603 | cdcfcaf | short |
| 873.msn_604 | cdcfcaf | short |
| 873.msn_605 | cdcfcaf | short |
| 873.msn_606 | cdcfcaf | short |
| 873.msn_607 | cdcfcaf | short |
| 873.msn_608 | cdcfcaf | short |
| 873.msn_609 | cdcfcaf | short |
| 873.msn_610 | cdcfcaf | short |
| 873.msn_end | cdcfcaf | mailer |

FBSNG on the web
Farm: CAF
Time: Thu May 23 01:48:13 2002
Report: Section 873.msn_600 status

Queues Jobs Nodes Process Types

User Monitor

ID: 873.msn_600 User: cdcfcf

Queue: msn Process Type: short

NProc: 1 Status: running

Need: 0 Depends:

Submitted: 05/23 01:46:57 Started: 05/23 01:47:09

CPU time limit: 2h00m

Proc Rsrc: cpu:100 disk:15 Sect Rsrc:

Command: /fbsng/caflcal/v1.01/CafExe cdcfcf@fcdhead1.fnal.gov/home/cdcfcf/v1.01/submitter/cafln/msn_%s.tgz msn@fcdflnx2.fnal.gov/cdf/scratch/msn/temp600.tgz msn 4h
cdcfcaf@fcdhead1.fnal.gov/home/cdcfcf/v1.01/submitter/fbs/FBS_%s.msn_600.1.log ./simple.sh 600

Other sections: msn_600 (running) msn_601 (running) msn_602 (running) msn_603 (running) msn_604 (running) msn_605 (running) msn_606 (running) msn_607 (running) msn_608 (running) msn_609 (running) msn_610 (running) msn_end (waiting)

Processes

| Process # | Node | Status | CPU Time | PID | Command |
|-----------|-------------|---------|----------|------|---|
| 1 | fcdcfcaf057 | running | 0 | 6931 | CafExe cdcfcf@fcdhead1.fnal.gov/home/cdcfcf/v1.01/submitter/cafln/msn_%s.tgz msn@fcdflnx2.fnal.gov/cdf/scratch/msn/temp600.tgz msn 4h |
| | | | 0 | 6940 | simple.sh 600 |
| | | | 0 | 7221 | sleep 120 |

FCS Group | FBSNG

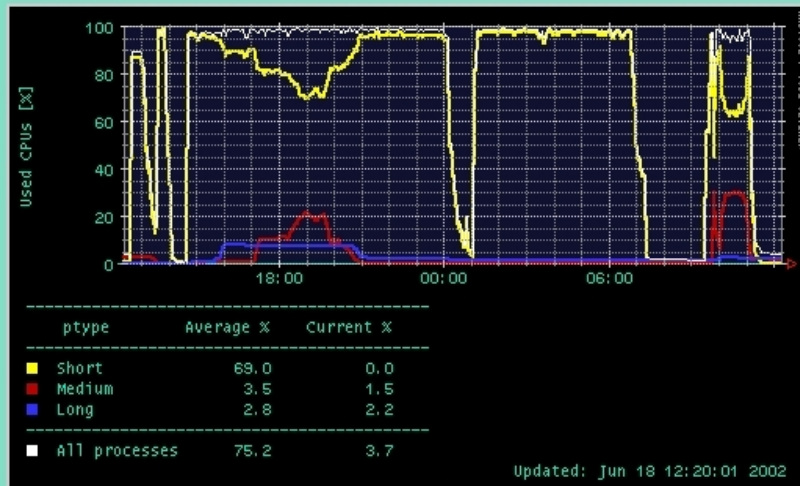
FBSWWW version 0.1



CAF Utilization



CAF global status



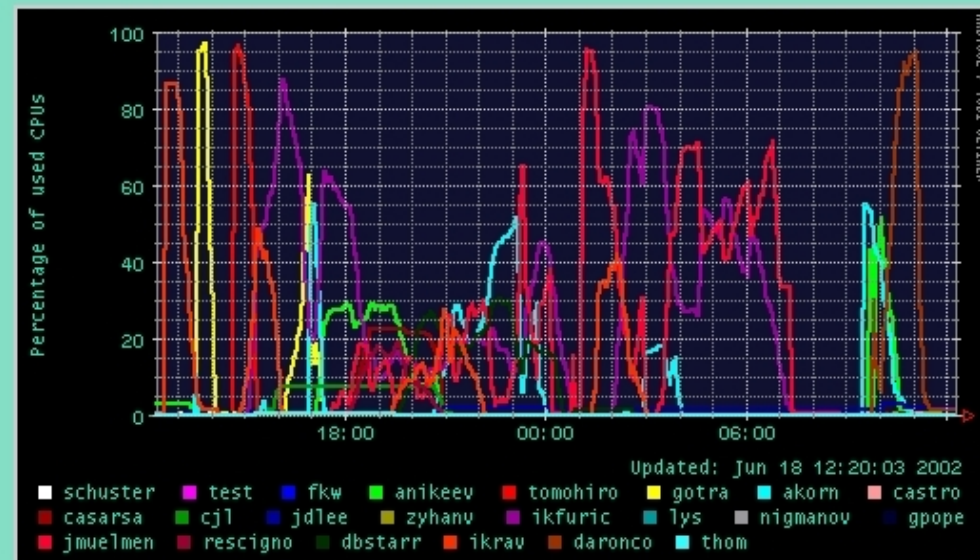
Built using RRDTool

Status summary

| | Short | Medium | Long | All Types |
|-------------------------------------|-------|--------|------|-----------|
| Running sections | 0 | 2 | 3 | 5 |
| Pending sections | 0 | 0 | 0 | 0 |
| Waiting time [hh:mm] (24h average): | | | | |
| per job | 0:04 | 0:26 | 0:00 | 0:15 |
| per section | 2:12 | 0:52 | 0:00 | 1:32 |
| Running time [hh:mm] (24h average) | 0:20 | 0:35 | 0:00 | 0:27 |

Updated: Jun 18 12:20:02 2002

Active queues (last 24h)



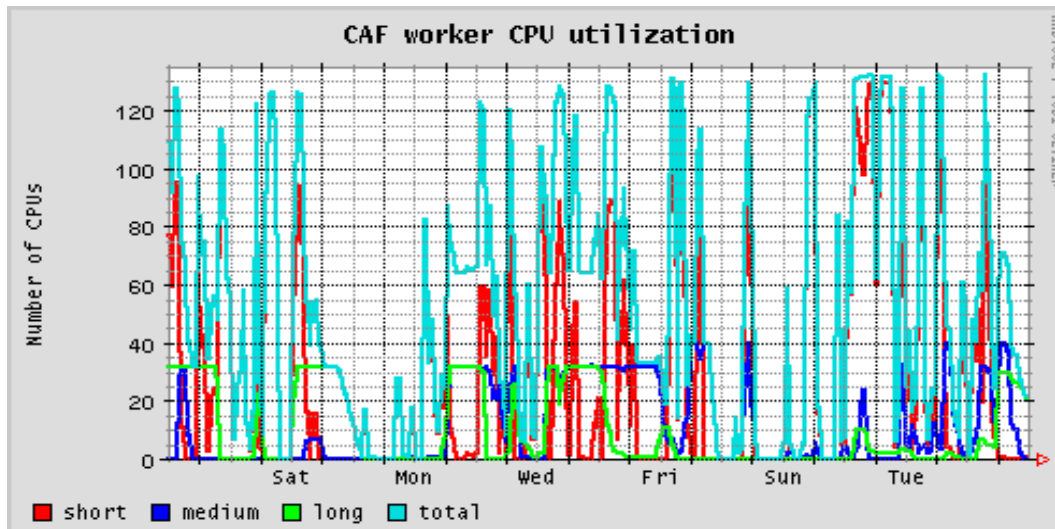
Built using RRDTool

CAF in active use by CDF collaboration

- 120 CAF Users (queues) to date
- 2-5 new users per day
- Several dozen simultaneous users in a typical 24 hr period

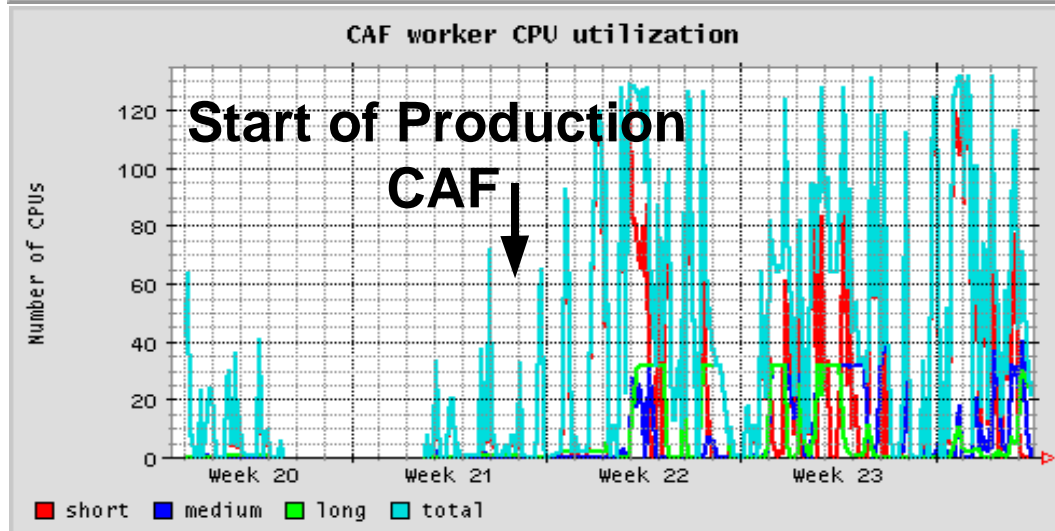


CAF Utilization



CAF utilization steadily rising since opened to collaboration

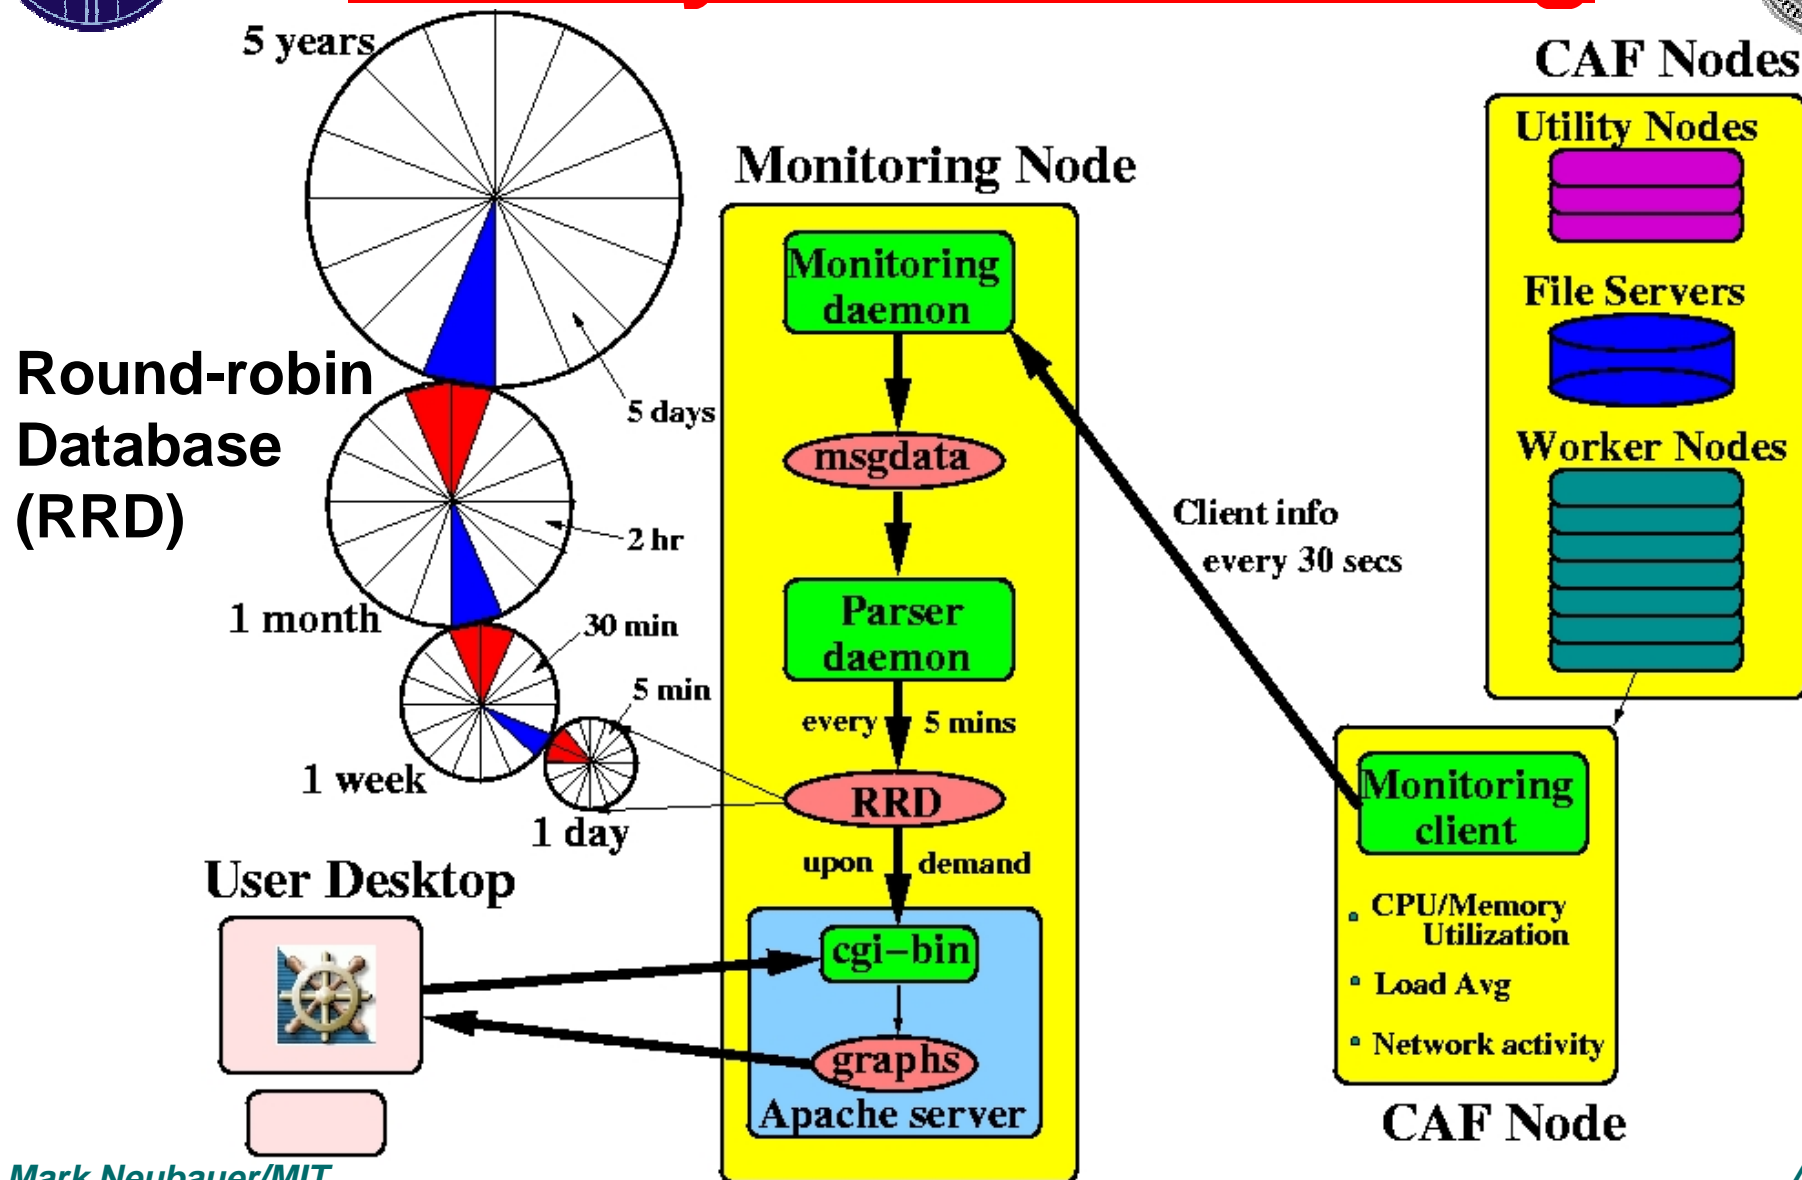
Providing 10-fold increase in analysis resources for summer physics conferences



Need for more CPU on the horizon



CAF System Monitoring



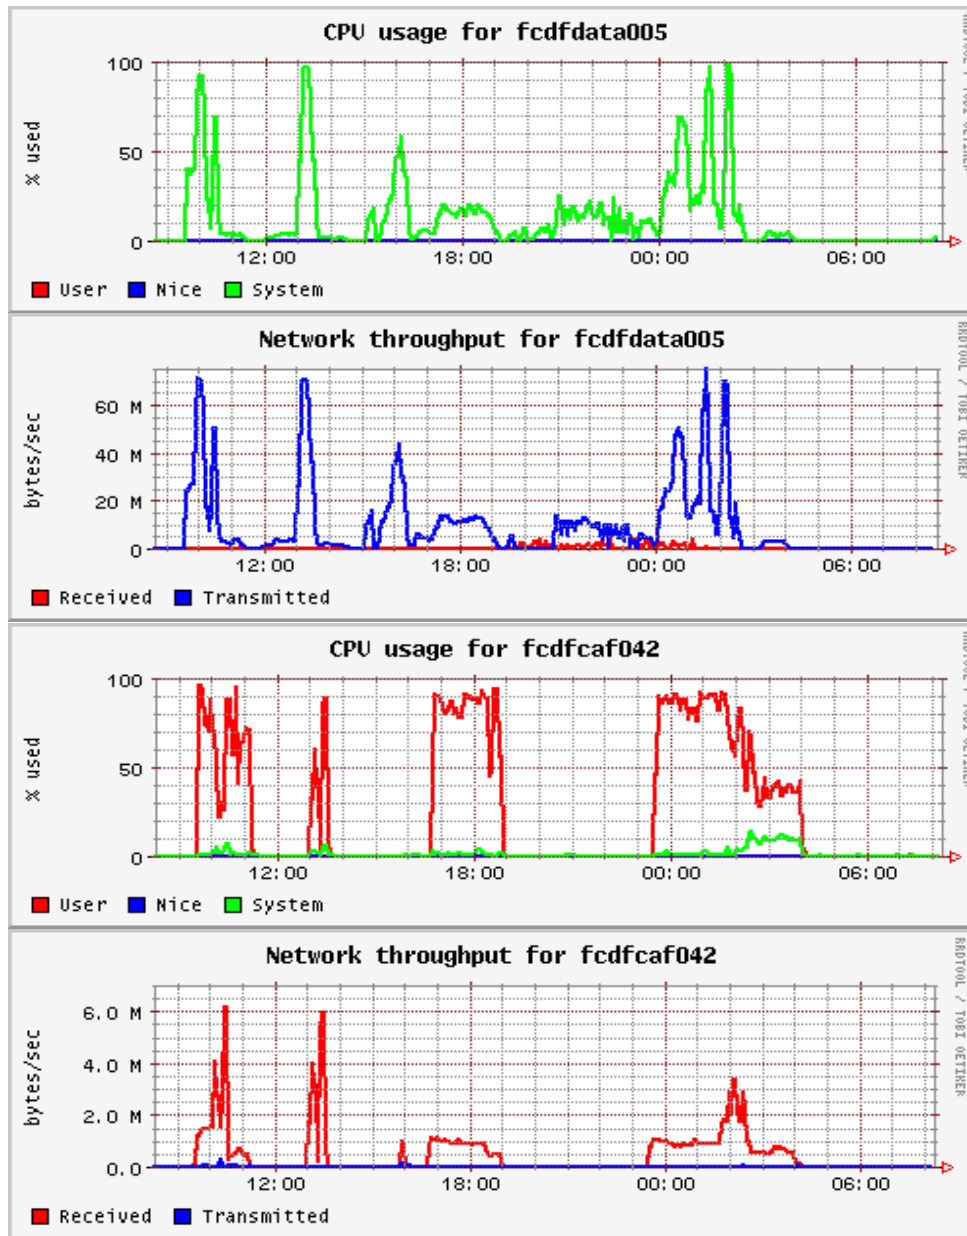
System Monitoring

2 TB File Server

Data transfers CPU limited

Analysis Jobs CPU bound

Worker Node



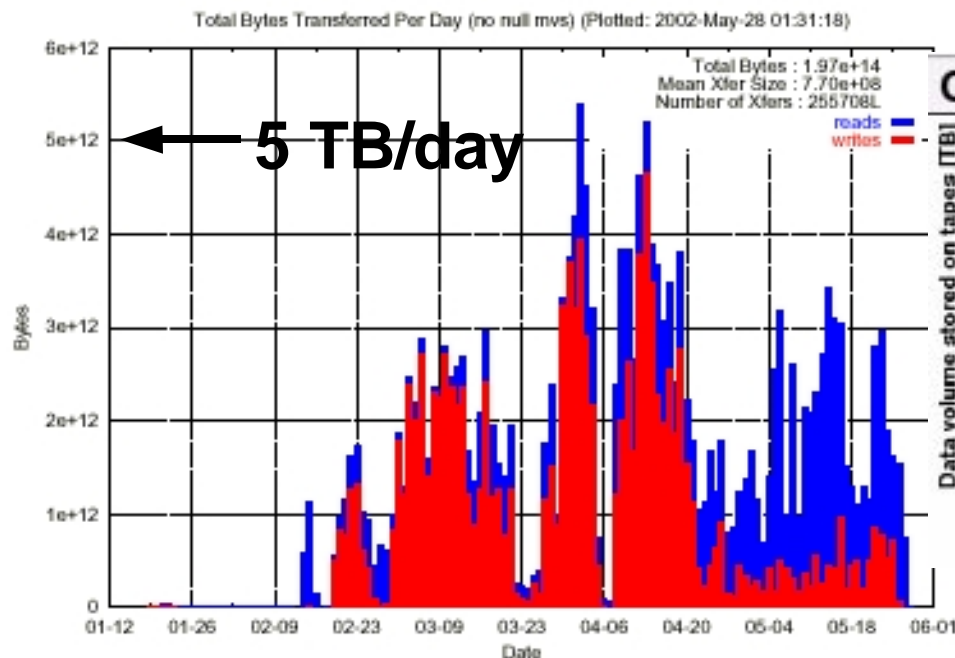


Data Handling

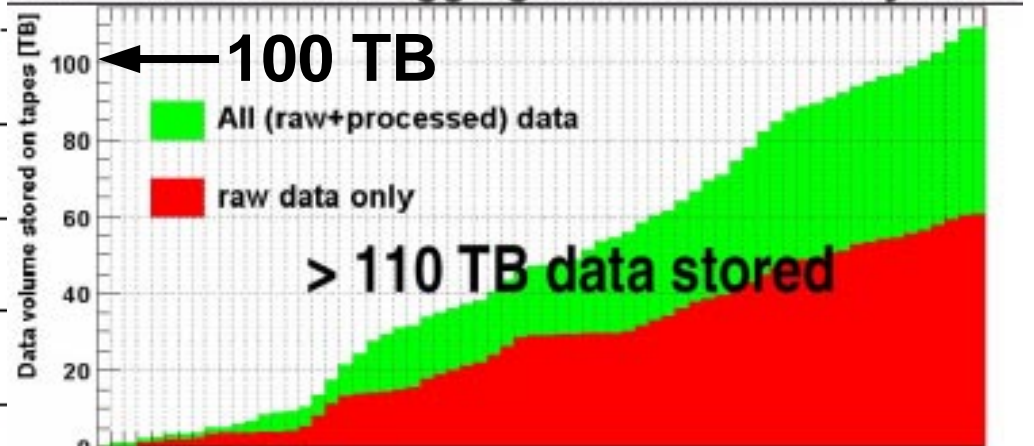


Data archived using STK 9940 drives and tape robot

Enstore: Network-attached tape system developed at FNAL
→ provides interface layer for staging data from tape



CDF Run II Data Logging March 2001 - May 2002





Data Handling



Dcache → network-attached disk cache from DESY

- Front-end disk cache for Enstore (read and write disk pools)
- Currently in β testing → working toward production use in CDF

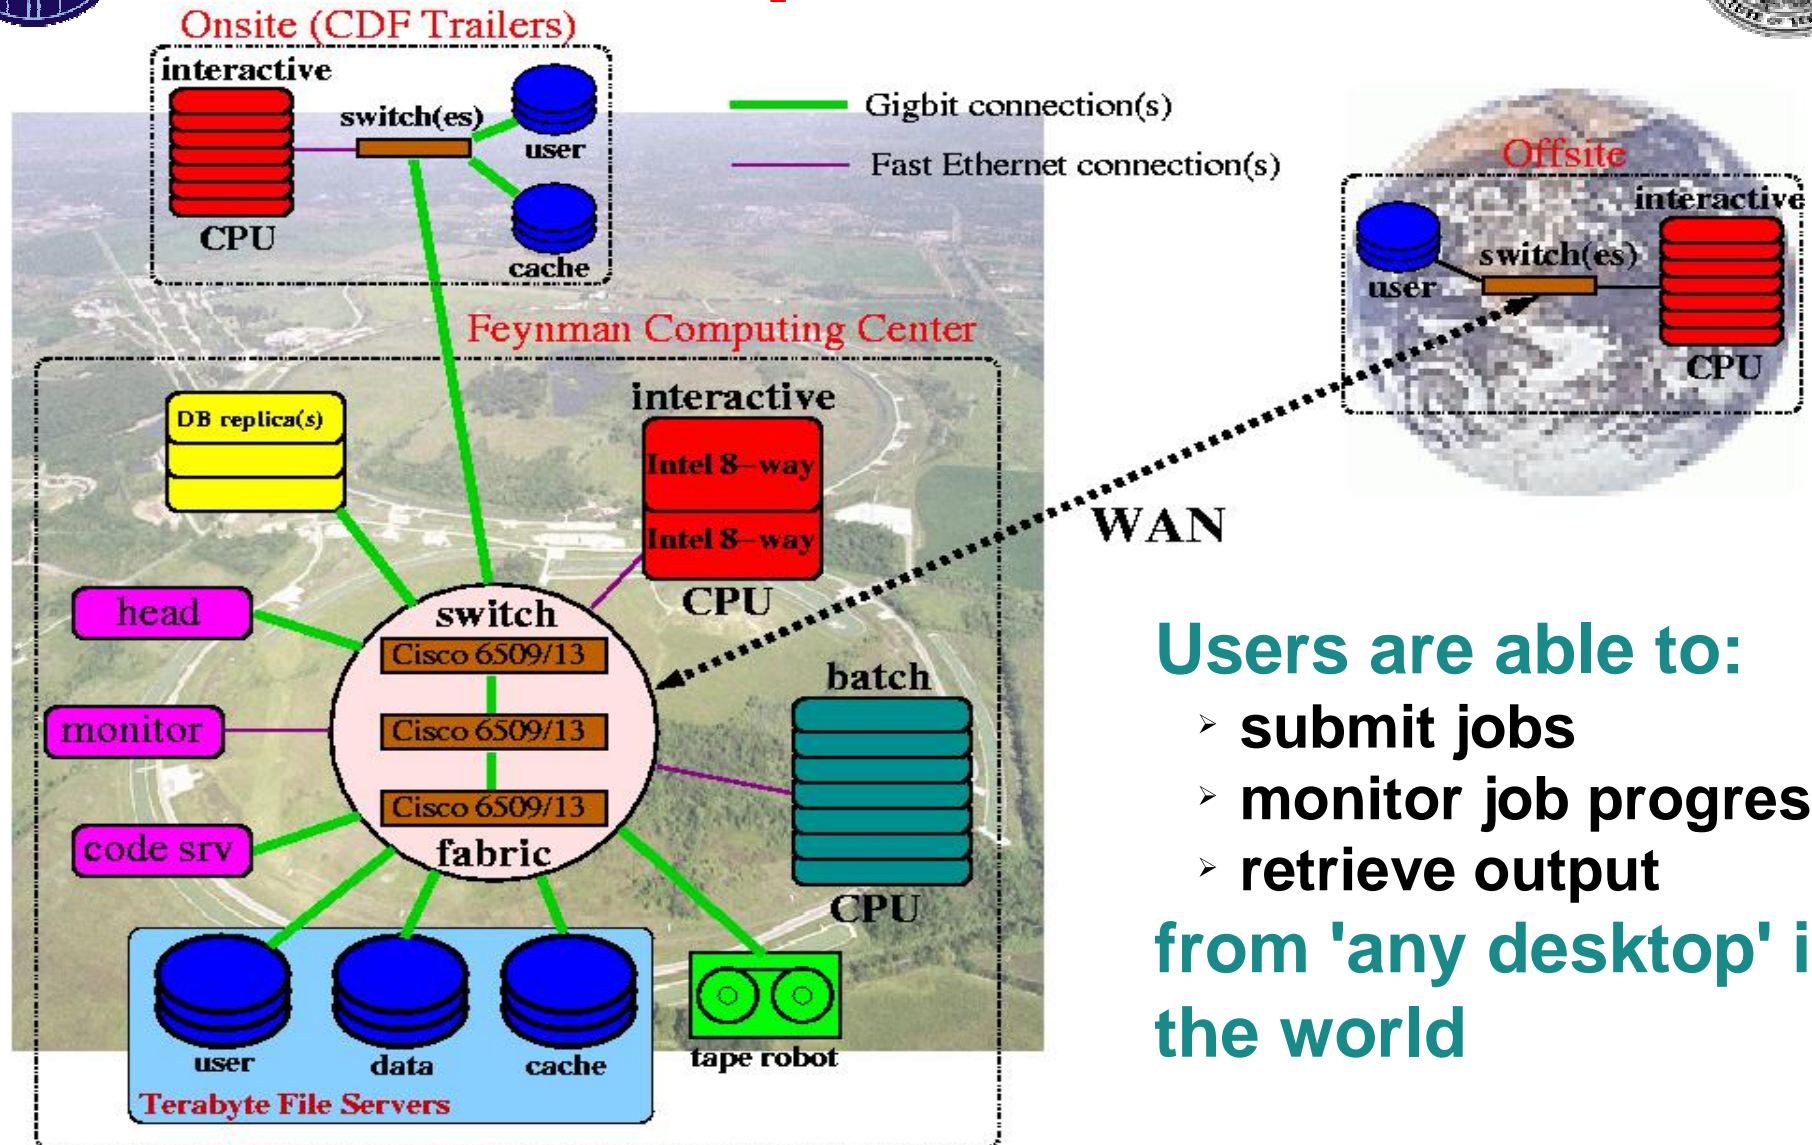
SAM → framework for global data handling/distribution

- Jointly developed by FNAL Computing division and D0
- Works with Enstore and CDF analysis software framework
- Currently under evaluation for use in CDF data distribution

→ see Igor Terekhov's talk



CAF Implementation



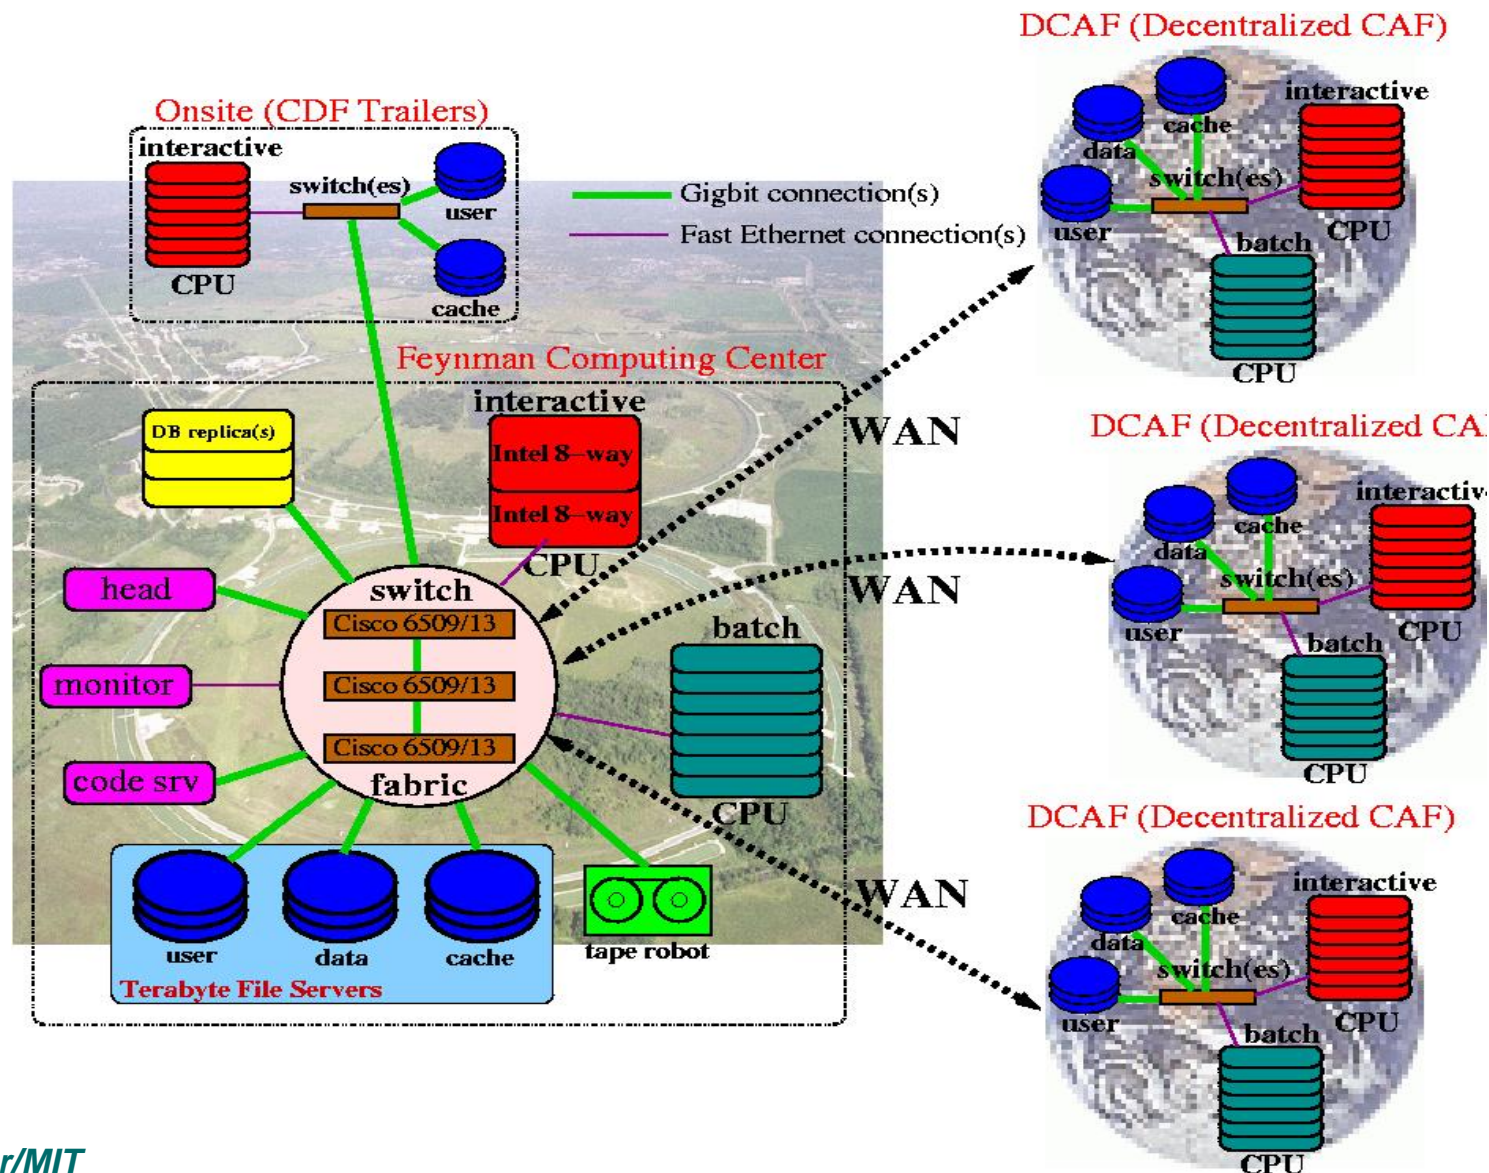
Users are able to:

- submit jobs
- monitor job progress
- retrieve output

from 'any desktop' in the world

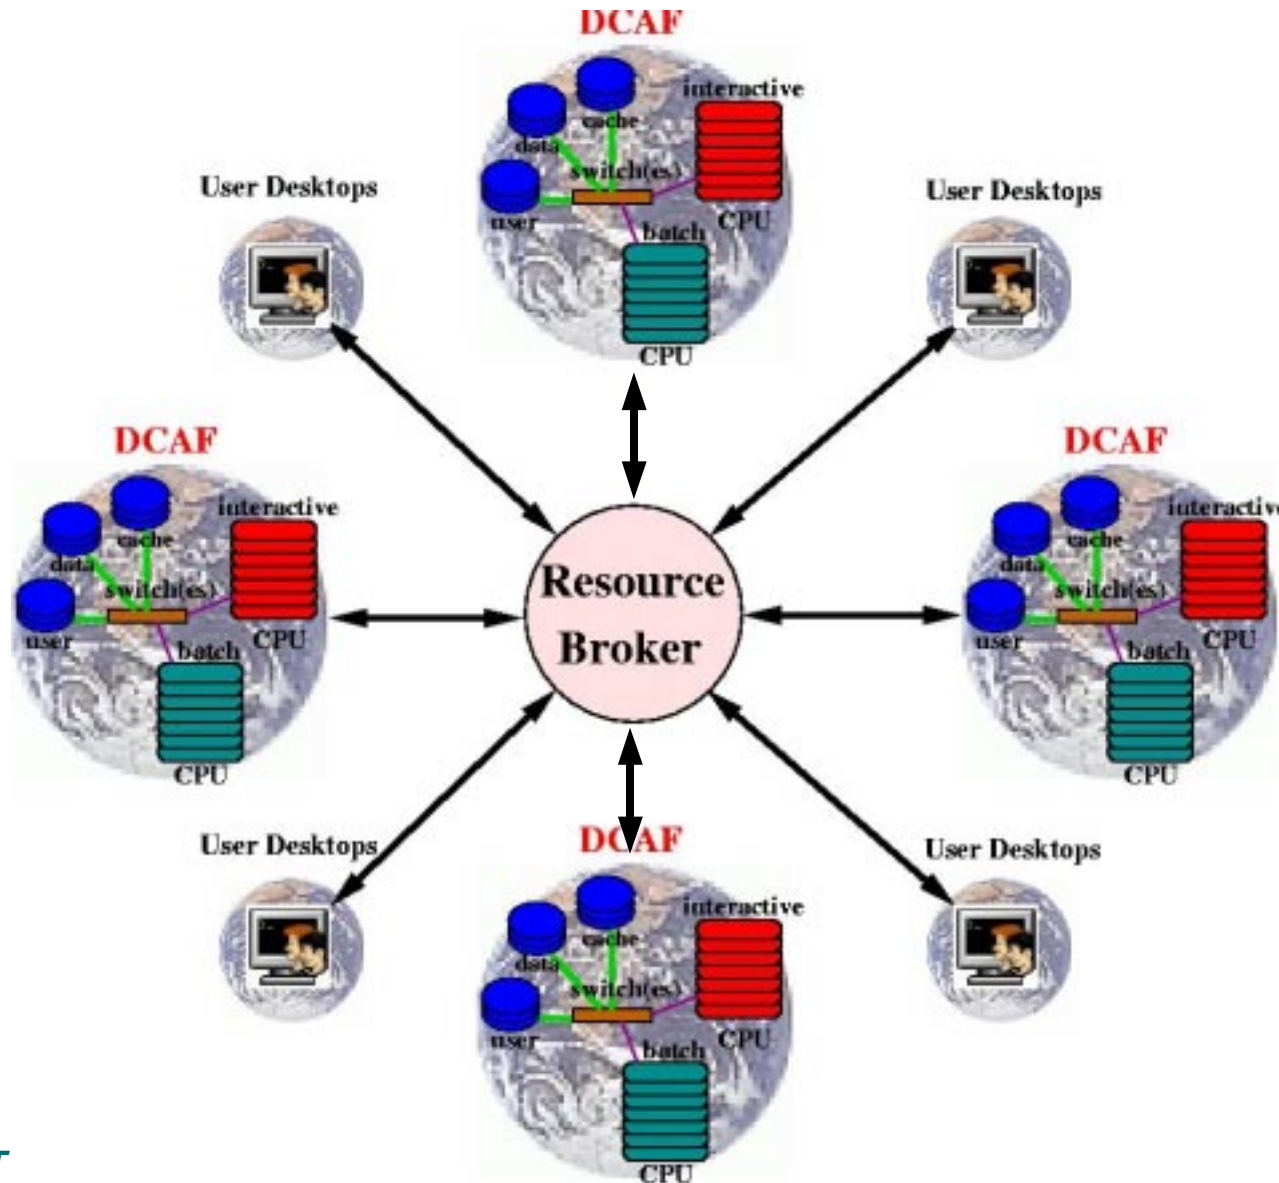


Toward the Grid





Peer-to-Farm Paradigm





Brokering scheme



Minimize job execution time:

DCAFs update broker

- CPU/disk utilization
- Local data

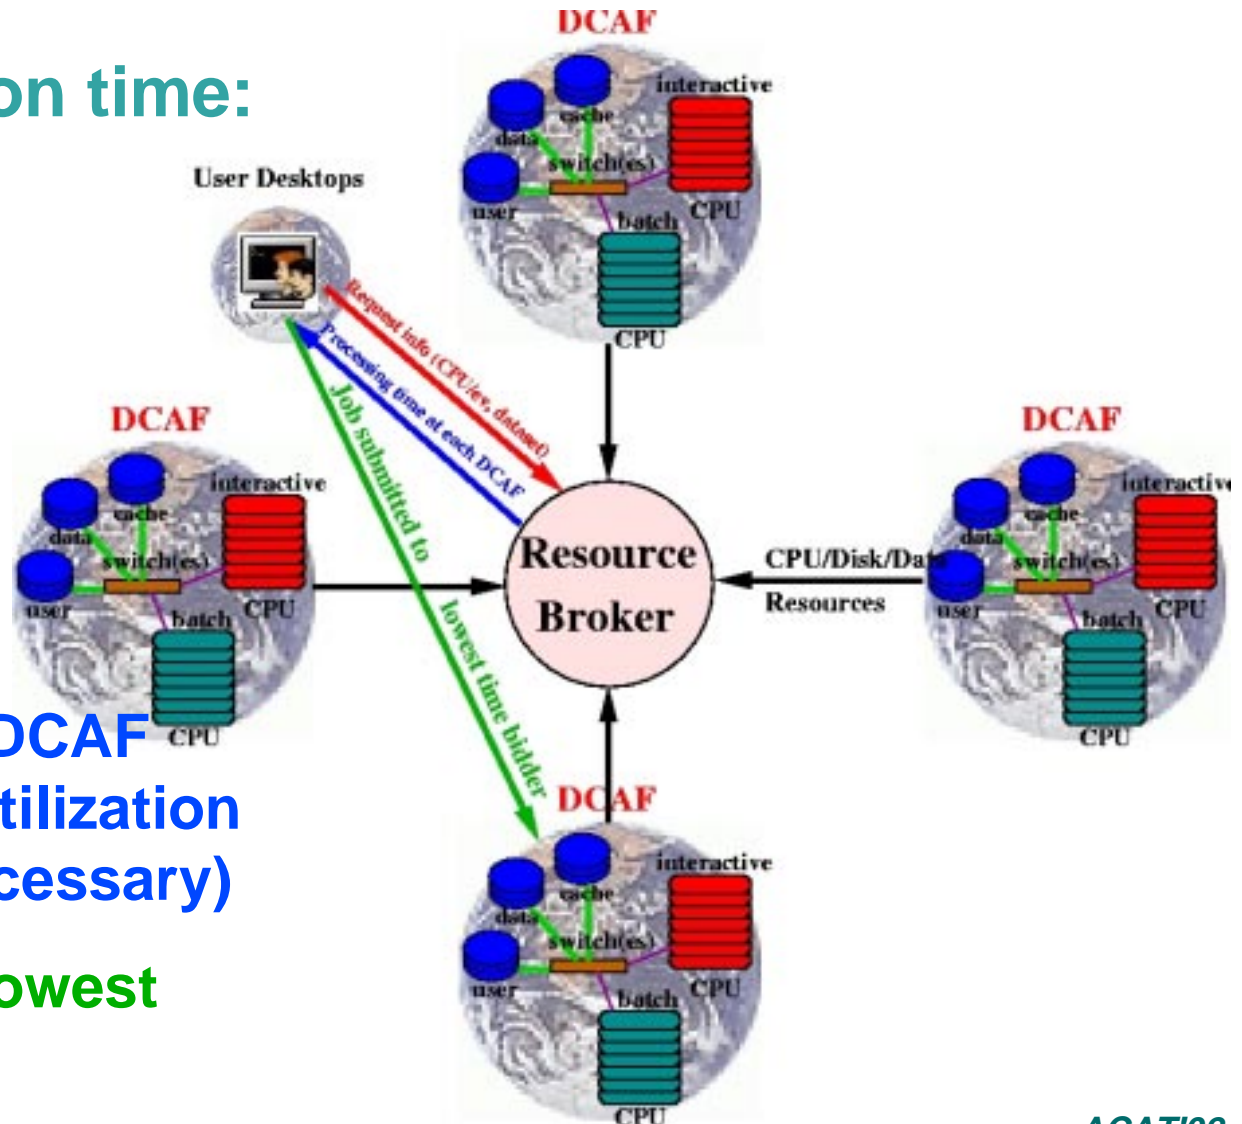
User generates request

- CPU time/event
- Metadata ID (dataset)

Execution time on each DCAF

- CPU+I/O resources+utilization
- Data movement (if necessary)

Job goes to DCAF with lowest 'bid'





Summary/Conclusions



Distributed Peer-to-Farm Computing Model

Production system under heavy use:

- Single farm at FNAL
- Many peers all over the world
 - 100+ total users
 - 100+ simultaneous jobs
 - Regularly up to 800 jobs per user queued

Future development:

- Extend data handling
- Multi-farm brokering
- Scale system by $O(10)$